Supplementary

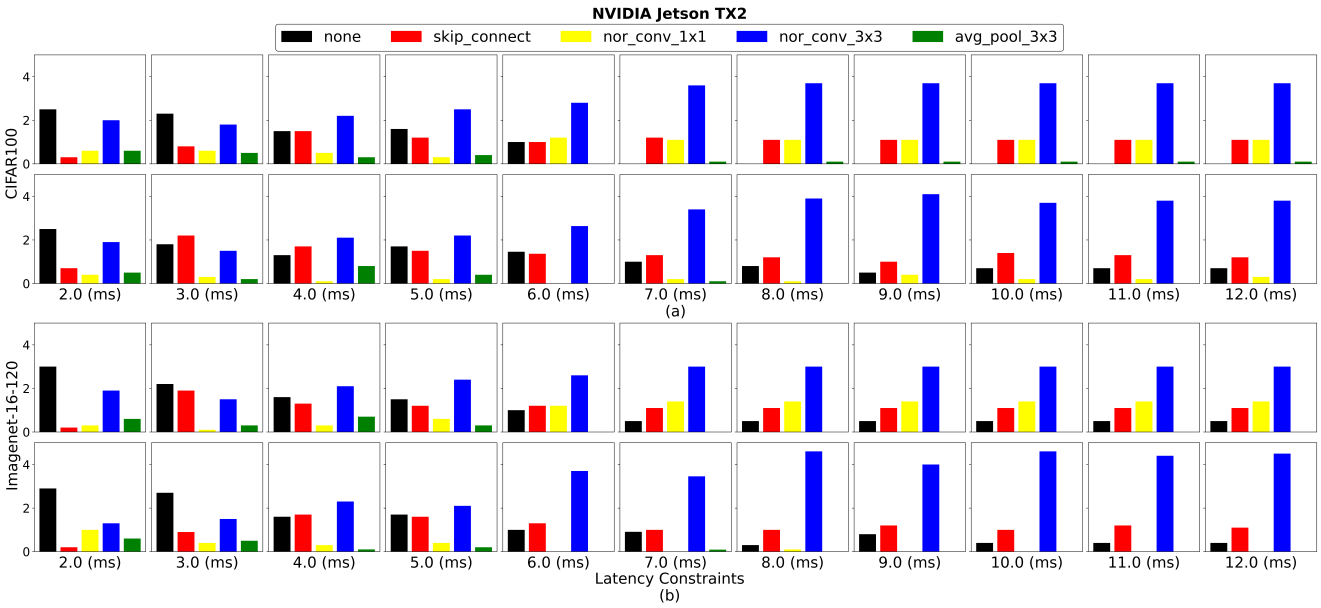

Figure 1: Results of HW-NAS with different latency constraints for (a) CIFAR100 (b) ImageNet16-120 experiments. The top row of each experiment shows the average number of different types of operations (i.e. convolution 3x3, convolution 1x1, max pooling 3x3, skip connect and none) present in the top-10 architectures of NAS-Bench-201 in terms of test accuracy. The bottom row of each experiment shows the same quantities present in the architectures discovered from ten independent architecture search runs. Note that the architecture search is performed in the NAS-Bench-201 search space for NVIDIA Jetson TX2 as the target hardware in this case. The similarity of the distribution of operations at different latency constraints may suggests that the proposed method converges to the high performing architecture sub-space.

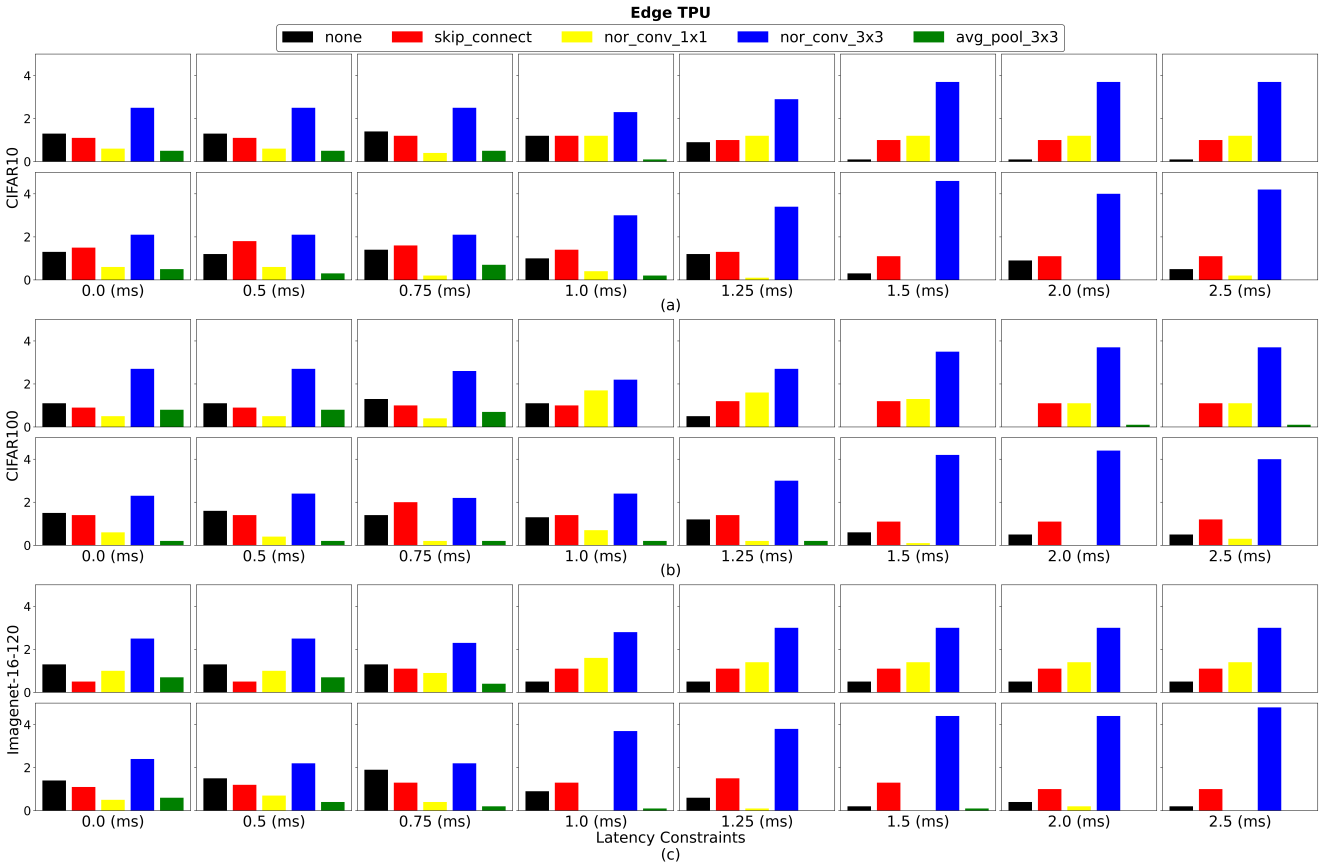

Figure 2: Same experiments as in Figure 1 with different latency constraints for Edge TPU on (a) CIFAR10 (b) CIFAR100 (c) ImageNet16-120.

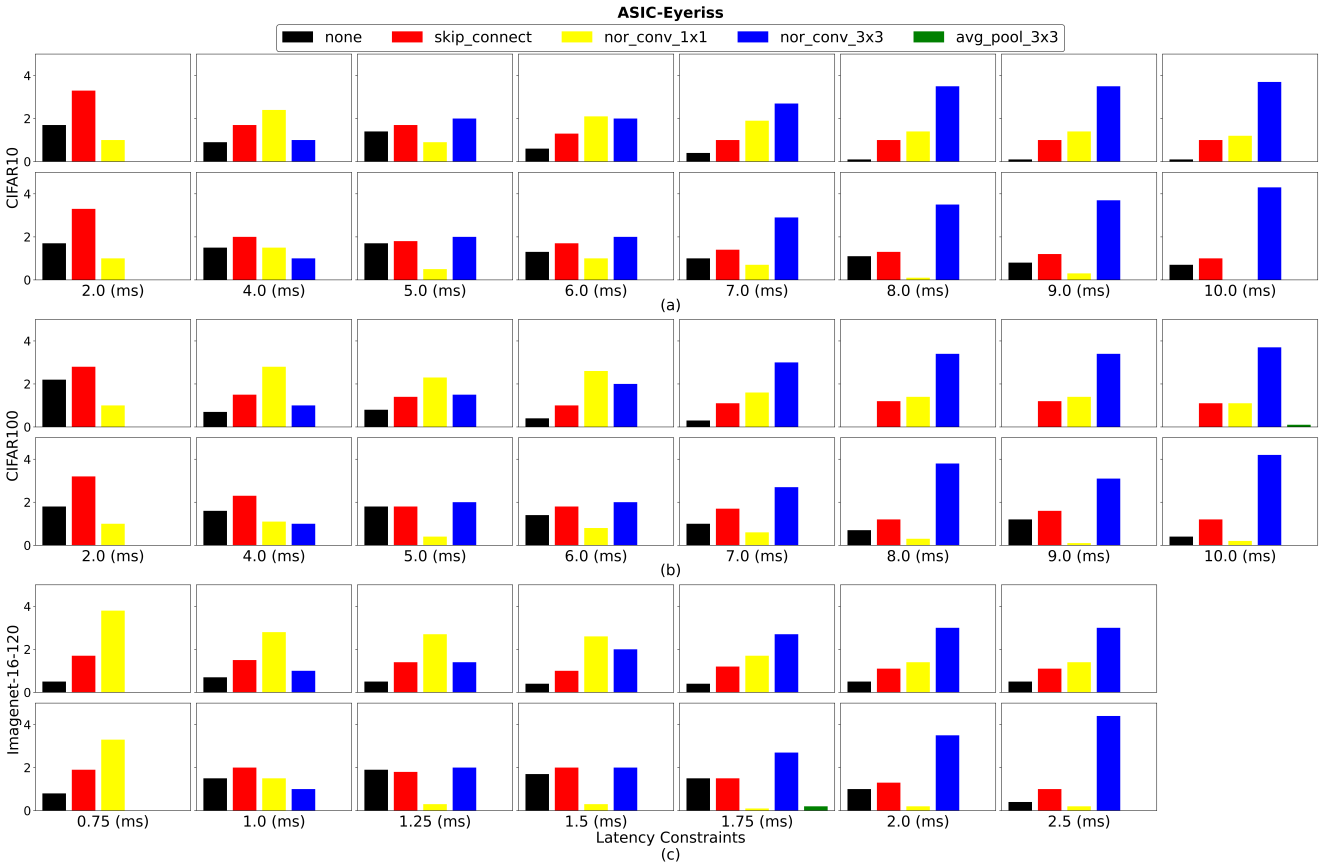

Figure 3: Same experiments as in Figure 1 with different latency constraints for ASIC-Eyeriss on (a) CIFAR10 (b) CIFAR100 (c) ImageNet16-120.

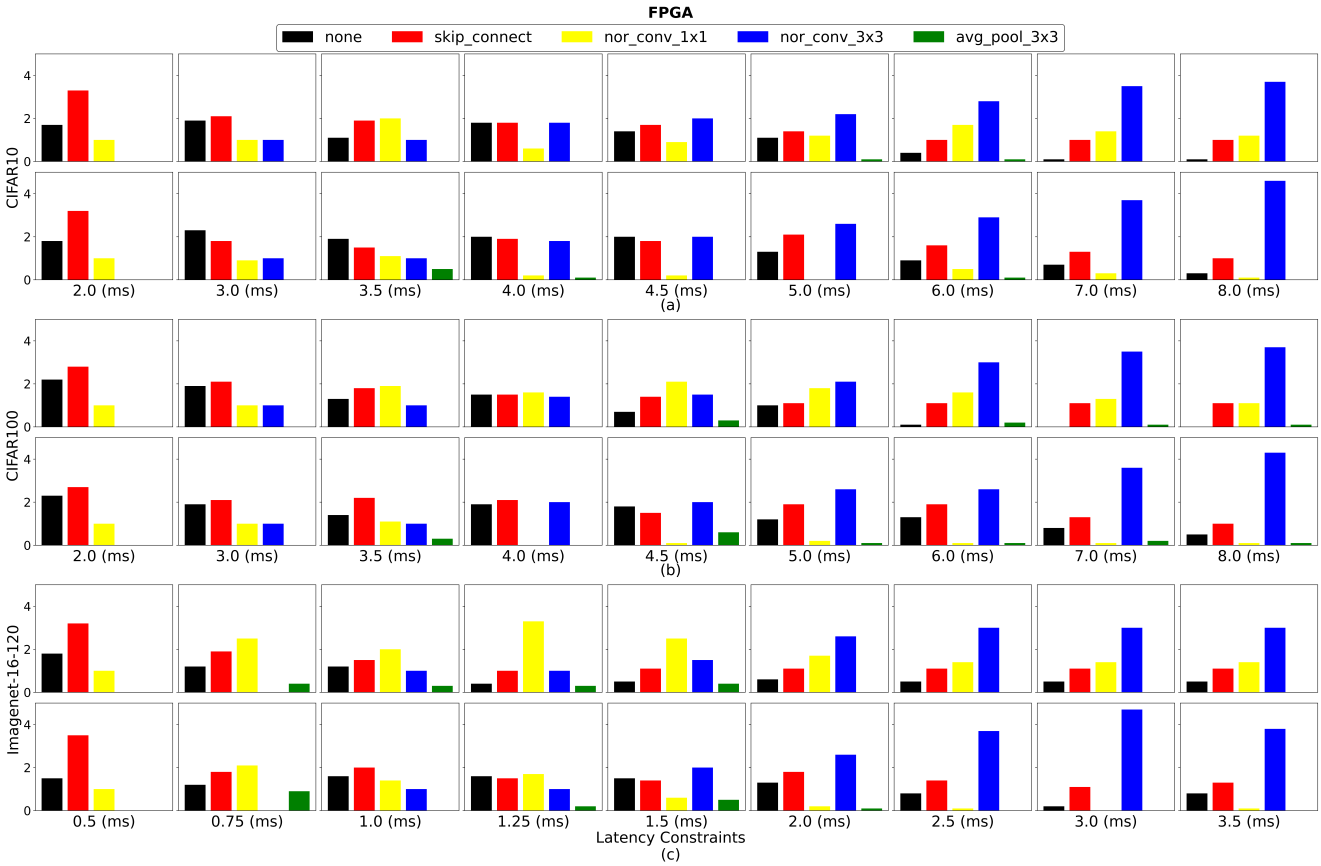

Figure 4: Same experiments as in Figure 1 with different latency constraints for FPGA on (a) CIFAR10 (b) CIFAR100 (c) ImageNet16-120.

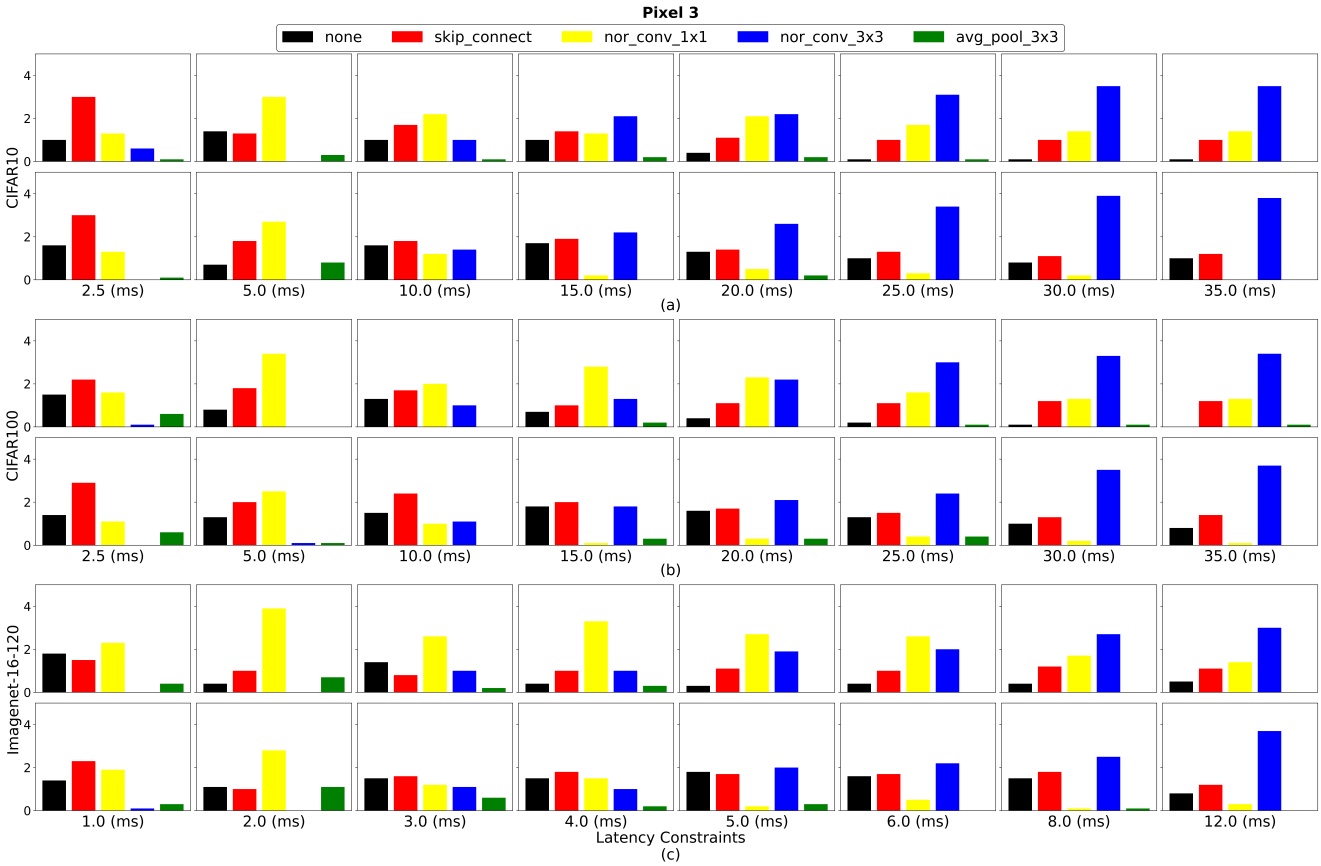

Figure 5: Same experiments as in Figure 1 with different latency constraints for Pixel 3 on (a) CIFAR10 (b) CIFAR100 (c) ImageNet16-120.

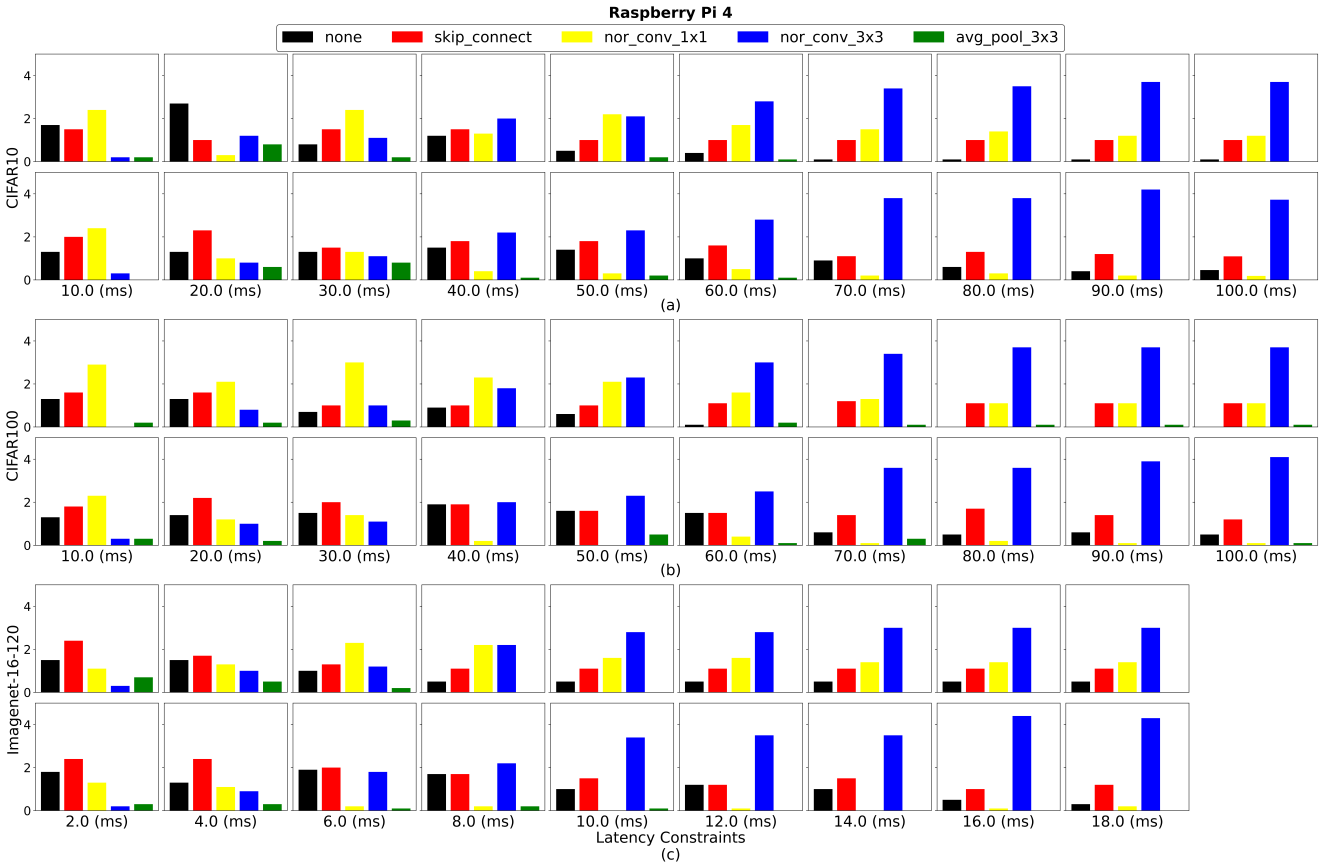

Figure 6: Same experiments as in Figure 1 with different latency constraints for Raspberry Pi 4 on (a) CIFAR10 (b) CIFAR100 (c) ImageNet16-120.

1. Datasets

CIFAR-10 and **CIFAR-100** [3] have 50K training images and 10K testing images with images classified into 10 and 100 classes respectively. **ImageNet**[2] is a well known benchmark for image classification containing 1K classes with 1.28 million training images and 50K test images. **ImageNet-16-120** [1] is a variant of ImageNet which is downsampled to 16x16 pixels with labels $\in [0, 120]$ to construct ImageNet-16-120 dataset.

References

[1] Patryk Chrabaszcz, Ilya Loshchilov, and Frank Hutter. A downsampled variant of imagenet as an alternative to the cifar datasets. *arXiv preprint arXiv:1707.08819*, 2017.

[2] J. Deng, W. Dong, R. Socher, L. Li, Kai Li, and Fei-Fei Li. Imagenet: a large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009.

[3] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.